

CLAIMS

What is claimed is:

- 1 1. A method for accessing data records from a large set of records stored in a
2 database, comprising:
 - 3 defining a plurality of boundaries to logically partition the large set of records
 - 4 into a plurality of buckets in accordance with a predefined sort order, said sort order
 - 5 based on data stored in at least two columns of a database table in which at least a
 - 6 base portion of each data record is stored;
 - 7 in response to a request to retrieve a data record,
 - 8 determining a bucket in which the data record is contained; and
 - 9 retrieving a subset of the large set of records from the database
- 10 corresponding to data records contained in the bucket.

- 1 2. The method of claim 1, wherein retrieving the subset of the large set of
2 records comprises:
 - 3 searching the plurality of boundaries for a sequentially-adjacent pair of
 - 4 boundaries as defined by the predetermined sort order that bound the data record in
 - 5 the request;
 - 6 generating a query including a search clause defining a data set that is
 - 7 bounded by the sequentially-adjacent pair of boundaries; and
 - 8 executing the query against the database.

1 3. The method of claim 1, further comprising storing a boundary marker for each
2 boundary comprising data pertaining to said at least two columns of the database
3 table.

1 4. The method of claim 3, wherein each boundary marker comprises first,
2 second and tertiary data values pertaining to data values stored in respective
3 columns in the database table, said data values pertaining to an actual record or
4 pseudo record to which that boundary marker corresponds.

1 5. The method of claim 3, wherein the boundary markers are stored as a
2 concatenated list.

1 6. The method of claim 3, wherein the boundary markers are stored as a set of
2 respective records in the database.

1 7. The method of claim 1, wherein the subset of data records corresponding to
2 the bucket is provided to a batch processing component, further comprising
3 retrieving other buckets of data and providing them to the batch processing
4 component until data records corresponding to the entire large set of data records
5 have been provided to the batch processing component.

1 8. The method of claim 1, further comprising:
2 monitoring user navigation events;
3 sending data corresponding to a current viewset to a client, said viewset
4 comprising a subset of data records in the data bucket and corresponding to a
5 current user navigation context; and

6 in response to a user navigation event corresponding to a request to navigate
7 to a new navigation context, updating the current viewset with a new viewset by
8 sending a new subset of data records in the data bucket, wherein the new viewset
9 corresponds to the new navigation context.

1 9. The method of claim 8, further comprising:
2 determining if a user navigation event corresponds to a request to view one
3 or more data records that are not in a current data bucket;
4 and, in response to such a determination,
5 formulating a query request corresponding to a new data bucket that contains
6 the data records that are requested to be viewed;
7 executing the query to retrieve data records corresponding to the new data
8 bucket from the database; and
9 providing data corresponding to a new viewset comprising a subset of the
10 data records in the new data bucket to the client.

1 10. The method of claim 8, wherein the client is a web browser and the data
2 corresponding to the viewsets are sent as HTML data over a network to a client
3 machine on which the web browser is running.

1 11. A method for accessing data records from a database, comprising:
2 determining a sort order in which the data records are to be provided, said
3 sort order including data from at least two columns of a database table from which
4 the data records are based;

5 defining a plurality of boundaries, each boundary corresponding to a relative
6 position of an actual or pseudo data record within a set of data records sorted in
7 accordance with the sort order;

8 in response to a request for a data record or a group of related data records
9 within the set of data records,

10 determining a lower and upper boundary corresponding to a data bucket
11 containing the data record or at least a portion of the group of related data records,
12 said lower and upper boundaries being selected from among the plurality of
13 boundaries points; and

14 retrieving data records logically contained within the data bucket based on
15 the relative positions of the lower and upper bucket boundaries within the sorted set
16 of data records.

1 12 The method of claim 1, wherein each boundary is marked by a boundary
2 marker comprising a unique set of data corresponding to said at least two columns
3 in the database table;

1 13. The method of claim 12, further comprising storing the boundary markers in
2 correspondence with the sort order.

1 14. The method of claim 11, wherein defining the boundaries comprises:
2 opening a data set corresponding to the data records, sorted in accordance
3 with the sort order and containing data corresponding to at least said at least two
4 columns in the database table; and
5 looping through the data set, defining boundaries every N rows in the data
6 set.

1 15. The method of claim 14, wherein subsets of the data set are opened on a
2 bucket-wise basis using an estimated set of boundaries, further comprising defining
3 actual boundaries to replace the estimated boundaries by looping through each
4 subset and defining actual boundaries every N rows of that subset.

1 16. The method of claim 11, wherein determining the lower and upper
2 boundaries comprises:

3 searching through the plurality of boundaries to identify the boundary having
4 a relative position in the sorted order that is closest to the data record or a first
5 record in the group of related data records from among boundaries having relative
6 positions prior to the data record or first record in the group of related data records
7 in the sorted order to determine the lower boundary; and

8 identifying the boundary having the next relative position after the boundary
9 determined for the lower boundary to determine the upper boundary.

1 17. The method of claim 11, further comprising:

2 monitoring a number of records returned by a query that is executed to
3 retrieve the data bucket; and

4 altering the query to dynamically adjust the size of subsequent buckets based
5 on the number of records returned by a prior query relative to a number of records
6 contained in a data bucket having a desired size.

1 18. The method of claim 17, wherein the size of the subsequent buckets are
2 adjusted by skipping a determined number of boundaries.

1 19. A machine-readable media on which a plurality of instructions are stored that
2 when executed by a computing machine perform the operations of:

3 receiving a request to retrieve data records from a database on which a large
4 set of records are stored;

5 searching a predefined set of boundaries that logically partition the large set
6 of records into a plurality of buckets in accordance with a predefined sort order that
7 is based on data stored in at least two columns of a database table in which at least
8 a base portion of each data record is stored;

9 determining a bucket in which the data record is contained;

10 formulating a query to retrieve a subset of the large set of records, said
11 subset corresponding to the data records contained in the bucket;

12 submitting the query to the database for execution; and

13 receiving the subset of records contained in the bucket from the database.

1 20. The machine-readable media of claim 19, wherein retrieving the subset of the
2 large set of records contained in the bucket comprises:

3 searching the plurality of boundaries for a sequentially-adjacent pair of
4 boundaries as defined by the predetermined sort order that bound the data record in
5 the request; and

6 formulating the query to include a search clause defining a data set that is
7 bounded by the sequentially-adjacent pair of boundaries.

1 21. The machine-readable media of claim 19, wherein execution of the
2 instructions further performs the operations of:

3 providing the subset of data records corresponding to the bucket to a batch
4 processing component running on the computing machine or on another computing
5 machine linked in communication with the machine;
6 iteratively formulating and submitting queries to the database to retrieve ,
7 other buckets of data and providing them to the batch processing component until
8 data records corresponding to the entire large set of data records have been
9 provided to the batch processing component.

1 22. The machine-readable media of claim 19, wherein execution of the
2 instructions further performs the operations of:
3 monitoring user navigation events received from a client;
4 sending data corresponding to a current viewset to the client, said viewset
5 comprising a subset of data records in the data bucket and corresponding to a
6 current user navigation context; and
7 in response to a user navigation event corresponding to a request to navigate
8 to a new navigation context, updating the current viewset with a new viewset by
9 sending a new subset of data records in the data bucket, wherein the new viewset
10 corresponds to the new navigation context.

1 23. The machine-readable media of claim 22, wherein execution of the
2 instructions further performs the operations of:
3 determining if a user navigation event corresponds to a request to view one
4 or more data records that are not in a current data bucket;
5 and, in response to such a determination,
6 formulating a query request corresponding to a new data bucket that contains
7 the data records that are requested to be viewed;

8 executing the query to retrieve data records corresponding to the new data
9 bucket from the database; and
10 providing data corresponding to a new viewset comprising a subset of the
11 data records in the new data bucket to the client.

1 24. The machine-readable media of claim 22, wherein the client is a web browser
2 and the data corresponding to the viewsets are sent as HTML data over a network
3 to a client machine on which the web browser is running.

1 25. The machine-readable media of claim 19, wherein execution of the
2 instructions further performs the operations of:
3 monitoring a number of records returned by a query that is submitted to the
4 database to retrieve the bucket of data records; and
5 altering the query to dynamically adjust the size of subsequent buckets based
6 on the number of records returned by a prior query relative to a number of records
7 contained in a bucket having a desired size.

1 26. The machine-readable media of claim 25, wherein the size of the subsequent
2 buckets are adjusted by skipping a determined number of boundaries.

1 27. A machine-readable media on which a plurality of instructions are stored that
2 when executed by a computing machine perform the operations of:
3 defining a plurality of boundaries to logically partition a large set of records
4 into a plurality of buckets in accordance with a predefined sort order, said sort order
5 based on data stored in at least two columns of a database table in which at least a
6 base portion of each data record is stored;

7 storing a boundary marker for each boundary comprising a unique set of data
8 pertaining to said at least two columns of the database table.

1 28. The machine-readable media of claim 27, wherein each boundary marker
2 comprises first, second and tertiary data values pertaining to data values stored in
3 respective columns in the database table, said data values pertaining to an actual
4 record or pseudo record to which that boundary marker corresponds.

1 29. The machine-readable media of claim 27, wherein the boundary markers are
2 stored as a concatenated list.

1 30. The machine-readable media of claim 27, wherein the boundary markers are
2 stored as a set of respective records in the database..

1 31. The machine-readable media of claim 27, wherein the operation of defining
2 the boundaries comprises:

3 opening a data set corresponding to the large set of data records, sorted in
4 accordance with the sort order and containing data corresponding to at least said at
5 least two columns in the database table; and

6 looping through the data set, defining boundaries every N rows in the data
7 set.

1 32. The machine-readable media of claim 31, wherein subsets of the data set are
2 opened on a bucket-wise basis using an estimated set of boundaries, further
3 comprising defining actual boundaries to replace the estimated boundaries by

4 looping through each subset and defining actual boundaries every N rows of that
5 subset.

1 33. A computer system comprising:
2 a memory in which a plurality of machine instructions are stored;
3 a network interface to link the computer in communication with a database
4 server; and
5 a processor, coupled to the memory and the network interface, to execute the
6 plurality of machine instructions to cause the computer system to perform the
7 operations of:
8 receiving a request to retrieve data records from a database hosted by
9 the database server on which a large set of records are stored;
10 searching a predefined set of boundaries that logically partition the
11 large set of records into a plurality of buckets in accordance with a predefined
12 sort order that is based on data stored in at least two columns of a database
13 table in which at least a base portion of each data record is stored;
14 determining a bucket in which the data record is contained;
15 formulating a query to retrieve a subset of the large set of records, said
16 subset corresponding to the data records contained in the bucket;
17 submitting the query over the network interface to the database server
18 for execution; and
19 receiving the subset of records contained in the bucket from the
20 database server.

1 34. The computer system of claim 33, wherein retrieving the subset of the large
2 set of records contained in the bucket comprises:

3 searching the plurality of boundaries for a sequentially-adjacent pair of
4 boundaries as defined by the predetermined sort order that bound the data record in
5 the request; and

6 formulating the query to include a search clause defining a data set that is
7 bounded by the sequentially-adjacent pair of boundaries.

1 35. The computer system of claim 33, wherein execution of the machine
2 instructions further performs the operations of:

3 providing the subset of data records corresponding to the bucket to a batch
4 processing component running on the computing machine or on another computing
5 machine linked in communication with the machine;

6 iteratively formulating and submitting queries to the database to retrieve ,
7 other buckets of data and providing them to the batch processing component until
8 data records corresponding to the entire large set of data records have been
9 provided to the batch processing component.

1 36. The computer system of claim 33, wherein the computer system is linked in
2 communication with a client machine via the network interface and wherein
3 execution of the machine instructions further performs the operations of:

4 monitoring user navigation events received from the client machine;

5 sending data corresponding to a current viewset to the client machine, said
6 viewset comprising a subset of data records in the data bucket and corresponding to
7 a current user navigation context; and

8 in response to a user navigation event corresponding to a request to navigate
9 to a new navigation context, updating the current viewset with a new viewset by

10 sending a new subset of data records in the data bucket to the client machine,
11 wherein the new viewset corresponds to the new navigation context.

1 37. The computer system of claim 36, wherein execution of the machine
2 instructions further performs the operations of:

3 determining if a user navigation event corresponds to a request to view one
4 or more data records that are not in a current data bucket;

5 and, in response to such a determination,

6 formulating a query request corresponding to a new data bucket that contains
7 the data records that are requested to be viewed;

8 submitting the query request to the database server;

9 receiving corresponding to the new data bucket from the database server in
10 response to the query request; and

11 providing data corresponding to a new viewset comprising a subset of the
12 data records in the new data bucket to the client machine.

1 38. The computer system of claim 36, wherein a web browser is running on the
2 client machine and the data corresponding to the viewsets are sent as HTML data
3 via the network interface to the client machine.

1 39. The computer system of claim 33, wherein execution of the machine
2 instructions further performs the operations of:

3 monitoring a number of records returned by a query that is submitted to the
4 database server to retrieve the bucket of data records; and

5 altering the query to dynamically adjust the size of subsequent buckets based
6 on the number of records returned by a prior query relative to a number of records
7 contained in a bucket having a desired size.

1 40. The computer system of claim 39, wherein the size of the subsequent
2 buckets are adjusted by skipping a determined number of boundaries.

1

CONFIDENTIAL - ATTORNEY'S EYES ONLY